

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Utility Patent Application

SUSTAINABLE CAPACITY ESTIMATION

Inventor(s):

Laurent Massoulie

Peter Key

Dinan Gunawardena

CLIENT'S DOCKET NO. MS304969.2

ATTORNEY'S DOCKET NO. MS1-1635US

1

SUSTAINABLE CAPACITY ESTIMATION

2

3

Related Applications

4

5 This applications claims priority to U.S. Provisional Application
6 No. 60/474,656, entitled "Acappella - Admission Control and Active Probing for
7 Networked Dataflows", filed on May 30, 2003 and specifically incorporated
herein by referenced for all that it discloses and teaches.

8

Technical Field

9

10 The invention relates generally to networking environments, and more
11 particularly to estimation of sustainable capacity in a network.

12

Background

13

14 In a communication network, it can helpful to determine the network
15 capacity that is available for a new channel. For example, in a small home
16 network, a number of devices (e.g., servers, desktop computers, Personal Digital
17 Assistants (PDAs)) may be connected using a variety of OSI Communications
18 Model Layer 2 technologies. In addition, various combinations of switches, hubs,
19 wireless access points, and wireless bridges, etc. may be employed in the network
configuration.

20 Measuring network characteristics in such networks may be important for
21 applications and services such as streaming multimedia data across the network.
22 For example, with knowledge of the sustainable capacity between two nodes in the
23 network, a user's multimedia application can determine whether and at what data
24 rate to safely start a new multimedia stream.

However, existing approaches for determining available capacity generally require distributed intelligence in the network (e.g., detection of packet loss or management of Explicit Congestion Notification (ECN) marks). Furthermore, some such approaches are designed to operate only on wide area networks, failing to adequately support small network.

Summary

Implementations described and claimed herein address the foregoing problems by actively probing the network to estimate the sustainable capacity available in the network between two devices. One implementation probes the network using two probe sequences having different loads. Average round-trip delay times associated with each sequence are used as probe estimates to estimate the sustainable capacity in the network. Another implementation probes the network using two probe sequences, in which either the probe packet size or the inter-probe gap are different between the sequences. Calculation of a delay trend for each sequence leads to an estimate of the maximum network capacity and the background load on the network, from which sustainable capacity may be determined. Yet another implementation uses a windowing technique to probe the sustainable capacity in the network. The window size is increased until the network is saturated. The number of bytes transferred in the resulting window during a given time period is an estimate of the sustainable capacity of the network.

In some implementations, articles of manufacture are provided as computer program products. One implementation of a computer program product provides a computer program storage medium readable by a computer system and encoding a

1 computer program. Another implementation of a computer program product may
2 be provided in a computer data signal embodied in a carrier wave by a computing
3 system and encoding the computer program

4 The computer program product encodes a computer program for executing
5 on a computer system a computer process. A first probe estimate is computed
6 from a first probe sequence transmitted over a network from a first device to a
7 second device. A second probe estimate is computed from a second probe
8 sequence transmitted over the network from the first device to the second device.
9 The first probe sequence provides a different load to the network than the second
10 probe sequence. A sustainable capacity of the network is estimated based on the
11 first probe estimate and the second probe estimate.

12 In another implementation, a method is provided. A first probe estimate is
13 computed from a first probe sequence transmitted over a network from a first
14 device to a second device. A second probe estimate is computed from a second
15 probe sequence transmitted over the network from the first device to the second
16 device. The first probe sequence provides a different load to the network than the
17 second probe sequence. A sustainable capacity of the network is estimated based
18 on the first probe estimate and the second probe estimate.

19 In yet another implementation, a system is provided. A probe sequence
20 generator generates a first probe sequence and a second probe sequence. A
21 communication module transmits the first probe sequence and the second probe
22 sequence to a target. The communication module also receives a first response
23 sequence associated with the first probe sequence and a second response sequence
24 associated with the second probe sequence. The communications module
25 computes a first probe estimate from the first probe sequence and a second probe

1 estimate from the second probe sequence. A probe performance analyzer
2 estimates a sustainable capacity in the network characterizing sustainable capacity
3 of the network for data traffic from the first device and the second device based on
4 the first probe estimate and the second probe estimate.

5 Yet another computer program product implementation is provided. A first
6 round-trip delay trend associated with a first probe sequence transmitted over a
7 network from a first device to a second device is computed. A second round-trip
8 delay trend associated with a second probe sequence transmitted over the network
9 from the first device to the second device is computed. A maximum capacity and
10 a background load of the network are estimated based on the first round-trip delay
11 trend and the second round-trip delay trend. The sustainable capacity of the
12 network is estimated based on the maximum capacity and the background load of
13 the network.

14 Yet another method implementation is provided. A first round-trip delay
15 trend associated with a first probe sequence transmitted over a network from a first
16 device to a second device is computed. A second round-trip delay trend associated
17 with a second probe sequence transmitted over the network from the first device to
18 the second device is computed. A maximum capacity and a background load of
19 the network are estimated based on the first round-trip delay trend and the second
20 round-trip delay trend. The sustainable capacity of the network is estimated based
21 on the maximum capacity and the background load of the network.

22 Yet another system implementation is provided. A probe sequence
23 generator generates a first probe sequence and a second probe sequence. A
24 communication module transmits the first probe sequence and the second probe
25 sequence to a target. The communication module also receives a first response

1 sequence associated with the first probe sequence and a second response sequence
2 associated with the second probe sequence; and computes first round-trip delay
3 times of the first probe sequence and second round-trip delay times of the second
4 probe sequence. A probe performance analyzer computes a first round-trip delay
5 trend associated with a first probe sequence and a second round-trip delay trend
6 associated with a second probe sequence. The probe performance analyzer also
7 estimates a maximum capacity and a background load of the network based on the
8 first round-trip delay trend and the second round-trip delay trend, and estimates a
9 sustainable capacity of the network based on the maximum capacity and the
10 background load of the network.

11 Other implementations are also described and recited herein.

12 **Brief Descriptions of the Drawings**

13 FIG. 1 illustrates an exemplary network including a first network device
14 and a second network device in an asymmetric configuration.
15

16 FIG. 2 illustrates an exemplary sustainable capacity probe.
17

18 FIG. 3 illustrates a diagram depicting various elements of network capacity.
19

20 FIG. 4 illustrates an exemplary first probe sequence.
21

22 FIG. 5 illustrates an exemplary second probe sequence.
23

24 FIG. 6 illustrates an exemplary network including a first network device
25 and a second network device in a symmetric configuration.
26

27 FIG. 7 illustrates exemplary operations for estimating sustained capacity in
28 a network.
29

30 FIG. 8 illustrates alternative exemplary operations for estimating sustained
31 capacity in a network.
32

1 FIG. 9 illustrates exemplary windows for estimating sustained capacity in a
2 network.

3 FIG. 10 illustrates a system useful for implementing an embodiment of the
4 present invention.

Detailed Description

7 Various actively probing techniques on a network allow an estimation of
8 the sustainable capacity available in the network. One implementation probes the
9 network using two probe sequences having different loads. Average round-trip
10 delay times associated with each sequence are used as probe estimates to estimate
11 the sustainable capacity in the network. Another implementation probes the
12 network using two probe sequences, in which either the probe packet size or the
13 inter-probe gap are different between the sequences. Calculation of a delay trend
14 for each sequence leads to estimate of the maximum network capacity and the
15 background load on the network. Yet another implementation uses a windowing
16 technique to probe the sustainable capacity in the network. The window size is
17 increased until the network is saturated. The number of bytes transferred within
18 the window during a given time period provides an estimate of the sustainable
capacity of the network.

20 FIG. 1 illustrates an exemplary network 100 including a first network
21 device 102 and a second network device 104 in an asymmetric configuration. For
22 example, the first network device 102 may be a media server system and the
23 second network device 102 may be a client computer that has requested access to
24 streamed media. Other network devices 106 may also be coupled to the

1 network 100 and influence the sustainable capacity that is available for sending
2 data between devices 102 and 104.

3 The first network device 102 includes a transmission module, such as a
4 streaming media server 108, and a sustainable capacity probe 110. The streaming
5 media server 108 transfers data from the first network device 102 to the second
6 network device 104 through a logical network connection 112. In order to provide
7 a high quality of service to the second network device 104, the first network
8 device 102 attempts to estimate the sustainable capacity of the logical network
9 connection 112, which may also be heavily influenced by communications of
10 other network nodes, such as the devices 106.

11 The second network device 102 includes a reception module (or “IP
12 target”), such as a streaming media receiver 114. For example, the streaming
13 media server 108 may transmit streaming multimedia data over the logical
14 connection 112 to the streaming media receiver, which presents the multimedia
15 content to a user (e.g., through a media player module).

16 The logical network connection 112 includes a control channel 116 and a
17 data channel 118. The control channel and data channel may be unidirectional or
18 bidirectional. The control channel 116 communicates control signals used to
19 determine spare capacity between devices 102 and 104 (typically, bidirectional),
20 such as by communicating probe packets and acknowledgements. The probes and
21 acknowledgements may be communicated in-band or out-of-band with respect to
22 the data channel 118. The data channel 118 communicates the primary
23 transmission traffic, such as the streaming media. The transmission rate of this
24 multimedia data is dependent upon the sustainable capacity of the network 100
25 and the rate needed to serve the reception module with an acceptable quality of

1 service (e.g., the acceptable streaming rate). The control channel may be separate
2 from or integral with the application's data channel or the application's own
3 control channel.

4 To estimate the sustainable capacity in the network 100, the sustainable
5 capacity probe module 110 transmits probes to the second device 104. A probe
6 elicits a response from the second device 102. For example, a probe may take the
7 form of an ICMP ping packet, although other probe formats may be employed. In
8 another implementation, a probe may be included in an arbitrary IP packet. A
9 probe also includes a unique signature to identify the probe and the corresponding
10 response so that the sustainable capacity probe module 110 can associate the probe
11 and the response into a probe-response pair.

12 A probe and a response are also time-stamped. For example, the probe may
13 be time-stamped just prior to transmission and the response may be time-stamped
14 just after receipt to minimize unnecessary delay overhead between the two
15 communications. The sustainable capacity probe module 110 determines the
16 difference in the probe time stamp and the response time stamp of a probe-
17 response pair. In one or more embodiments, this difference value D may be used
18 to compute the estimated sustainable capacity of the network.

19 FIG. 2 illustrates an exemplary sustainable capacity probe module 200. By
20 way of a sustained capacity request 202, a transmission module inquires about
21 available sustained capacity in order to determine an appropriate transmission rate
22 from a first device to a second device. In one implementation, the sustained
23 capacity request 202 takes the form of a request for a specific level of sustained
24 capacity (e.g., "I want a usable capacity of X Bps for transmission to a device B ",
25 where device B is an IP target). However, other request formats are contemplated,

1 including without limitation a request regarding multiple acceptable usable
2 capacities or IP targets or a request for a best possible usable capacity or IP target.
3 Therefore, in various implementations, usable capacity requests may or may not
4 specify bounds on the requested capacities. The responses to such requests can
5 describe the feasibility of the capacity requests and may include an alternative set
6 of possible capacity bounds.

7 A probing module 204 operates in user mode on a first network device and
8 includes a probing sequence generator 206, which generates an appropriate probe
9 sequence that is transmitted to the IP target. However, it should be understood
10 that the probing module 204 could also operate in kernel mode or could even be
11 combined with application data. Various exemplary probe sequences are
12 discussed elsewhere in this document. Responsive to the sustained capacity
13 request 202, the probing module 204 identifies the IP target. Depending on the
14 probe sequence employed, the probing module 204 generates probes of a given
15 packet size and collects the probes into a sequence (or probe train) with a given
16 inter-probe gap T and a given number of probes N . In addition, in certain
17 implementations, an additional load may also be introduced into the probe
18 sequence.

19 A probe sequence is passed to a probe packet sender/receiver 206, which
20 (in the illustrated embodiment) executes in kernel mode; however, in alternative
21 implementations, the probe packet sender/receiver 206 may execute in a user
22 mode. The probe packet sender/receiver 206 adds a unique signature to outgoing
23 probe packets of a probe sequence. The probe packet sender/receiver 206 may
24 also be associated with the device's network stack, although this does not preclude
25 implementations of the probe packet sender/receiver 206 both within and outside

1 of the network stack. For example, in one implementation, the probing subsystem
2 is an operating system service that all applications can use. A probe packet
3 timestamp module 208 inserts a time stamp into each probe packet as it is
4 transmitted onto the network toward the IP target. The probe packet timestamp
5 module 208 may also be implemented in either a kernel mode or a user mode.

6 When each response packet (e.g., a ping acknowledgement) is received by
7 the probe packet timestamp module 208, the module timestamps the response
8 packet and checks the response packet to verify that the response corresponds to a
9 previously transmitted probe. The response packet is then passed to the probe
10 packet sender/receiver 206, which checks the unique signature of the response
11 packet and groups matching probe-response pairs.

12 Given a matching probe-response pair, the round-trip delay associated with
13 the pair is calculated and used by a probe performance analyzer 210 in the probing
14 module 204 to estimate the available sustained capacity for communications over
15 the network to the IP target. In one implementation, a sustained capacity
16 response 212 generated by the probe performance analyzer 210 specifies whether
17 a requested level of sustained capacity is available (e.g., YES/NO). Alternatively
18 or in addition, a sustained capacity response 212 may specify the estimated level
19 of sustained capacity available over the network to the IP target, or some other
20 combination of information responsive to the sustained capacity request 202.

21 FIG. 3 illustrates a diagram 300 depicting various elements of network
22 capacity. A capacity line 302 represents the available capacity of the channel
23 absent any traffic (i.e., the maximum network capacity). Of this maximum
24 network capacity, a portion of the traffic may be unreactive traffic (e.g., constant
25 bit rate (CBR) traffic), which is commonly communicated using Uniform

1 Datagram Protocol (UDP) packets, such as traffic associated with streaming media
2 traffic. In FIG. 3, this traffic is represented by region 304. In contrast, another
3 portion of the traffic may be reactive traffic, such as Transmission Control
4 Protocol (TCP) packets, as represented by region 306. The raw sustainable
5 capacity 308 is therefore represented by the maximum network capacity minus the
6 capacity occupied by the CBR and TCP traffic in the channel. In some cases, this
7 raw sustainable capacity is an acceptable value of the sustainable capacity in the
8 network.

9 However, reactive traffic (e.g., TCP traffic) may react to a loss of available
10 channel capacity by delaying or aborting transmissions. In the extreme, this
11 reaction may result in the removal of all reactive traffic from the channel. As
12 such, the maximum sustainable capacity 310 for streaming media traffic may be
13 represented by the maximum network capacity minus the CBR traffic. Applying a
14 load that saturates the network channel for a sustained period of time tends to
15 provide a more accurate estimate of the maximum sustainable capacity because
16 the reactive traffic is pushed out by the sustained, saturating load.

17 In one implementation, the sustainable capacity estimation involves two
18 probe sequences or probe trains. The network load introduced by each probe
19 sequence is different. For example, in FIG. 4, a low load probe sequence is
20 shown, without additional load packets; whereas in FIG. 5, a high load probe
21 sequence is shown, with additional load packets spaced between the probes.

22 FIG. 4 illustrates an exemplary first probe sequence 400 and a response
23 sequence 402. The sequences are laid out on time axes 404 and 406. Each probe
24 packet is sent periodically at time period T. The probe packets are designated in
25 FIG. 4 with an arrow pointing to the right and a capital letter. The capital letter

1 represents a unique signature inserted into the probe packet by the probe packet
2 sender/receiver. Each response packet is received after a round-trip to the IP
3 target, such as an ICMP ping ACK packet. The response packets are designated in
4 FIG. 4 with an arrow pointing to the left and a capital letter in italics. The capital
5 letter in the response packets associated the response with a specific probe packet.

6 Each probe packet is time-stamped before transmission, and each response
7 packet is time-stamped after reception. The difference in the reception timestamp
8 (e.g., as shown by vertical line 408) and the corresponding transmission timestamp
9 (e.g., as shown by vertical line 410) represents a round-trip delay time $D_{1,x}$, where
10 the subscript “1” represents a delay associated with a first probe sequence and the
11 subscript “x” represents the sequence number (e.g., 1, 2, 3,...) of the probe-
12 response pair. The average of the delays in the first probe sequence is indicated
13 as \hat{D}_1 .

14 FIG. 5 illustrates an exemplary second probe sequence 500 and a response
15 sequence 502. The sequences are laid out on time axes 504 and 506. Each probe
16 packet is sent periodically at time period T. In some implementations, T may
17 remain unchanged over long periods of communications. In alternative
18 implementations, T may be varied from one probe sequence to another in order to
19 negate correlation effects between probe traffic and background traffic.

20 The probe packets are designated in FIG. 5 with an arrow pointing to the
21 right and a capital letter. The capital letter represents a unique signature inserted
22 into the probe packet by the probe packet sender/receiver. Each response packet is
23 received after a round-trip to the IP target. The response packets are designated in
24 FIG. 5 with an arrow pointing to the left and a capital letter in italics. The
25

1 italicized capital letter in the illustrated response packets associates a response
2 with a specific probe packet.

3 In contrast to FIG. 4, the second probe sequence 500 includes load
4 packets 508. Although both the first and second probe sequences can include load
5 packets, the load introduced to the network channel should be different between
6 the two probe sequences. In one such implementation, any loads applied to not
7 saturate the network channel and, therefore, does not drive the network to a halt.

8 Each probe packet is time-stamped before transmission, and each response
9 packet is time-stamped after reception. The difference in the reception timestamp
10 (e.g., as shown by vertical line 510) and the corresponding transmission timestamp
11 (e.g., as shown by vertical line 512) represents a round-trip delay time $D_{2,x}$, where
12 the subscript “2” represents a delay associated with a second probe sequence and
13 the subscript “x” represents the sequence number (e.g., 1, 2, 3,...) of the probe-
14 response pair. The average of the delays in the second probe sequence is indicated
15 as \hat{D}_2 .

16 FIG. 6 illustrates an exemplary network including a first network
17 device 600 and a second network device 602 in a symmetric configuration. Other
18 network devices 606 may also be coupled to the network 100 and influence the
19 sustainable capacity that is available for sending data between devices 602
20 and 604.

21 The first network device 602 includes a transmission module, such as a
22 streaming media server 608, and a sustainable capacity probe 610. The streaming
23 media server 608 transfers data from the first network device 602 to the second
24 network device 604 through a logical network connection 612. In order to provide
25 a high quality of service to the second network device 604, the first network

1 device 602 attempts to estimate the sustainable capacity of the logical network
2 connection 612, which may also be heavily influenced by communications of
3 other network nodes, such as the devices 606. The second network device 602
4 includes a reception module (or “IP target”), such as a streaming media
5 receiver 614. For example, the streaming media server 608 may transmit
6 streaming multimedia data over the logical connection 612 to the streaming media
7 receiver, which presents the multimedia content to a user (e.g., through a media
8 player module).

9 The logical network connection 612 includes a control channel 616 and a
10 data channel 618. The data channel 618 may be bidirectional or there may be
11 separate server-client and client-server data and control channels. The control
12 channel 616 communicates control signals for probes and acknowledgements,
13 which may be communicated in-band or out-of-band with respect to the data
14 channel. The data channel 618 communicates the primary transmission traffic,
15 such as the streaming media. The transmission rate of this multimedia data is
16 dependent upon the sustainable capacity of the network 600 and the rate needed to
17 serve the reception module with an acceptable quality of service (e.g., the
18 acceptable streaming rate).

19 To estimate the sustainable capacity in the network 600 from the first
20 network device 602 to the second network device 604, the sustainable capacity
21 probe module 610 transmits probes to the second device 604. A probe elicits a
22 response from the second device 602. For example, a probe may take the form of
23 an ICMP ping packet, although other probe formats may be employed, generally
24 taking the form of IP packets. A probe also includes a unique signature to identify
25

1 the probe and the corresponding response so that the sustainable capacity probe
2 module 610 can associate the probe and the response of a probe-response pair.

3 A probe and a response are also time-stamped. For example, the probe may
4 be time-stamped just prior to transmission and the response is time-stamped just
5 after receipt to minimize unnecessary delay overhead between the two
6 communications. The sustainable capacity probe module 610 determines the
7 difference in the probe time stamp and the response time stamp of a probe-
8 response pair. This difference value D is used to compute the estimated
9 sustainable capacity of the network from the first network device 602 to the
10 second network device 604.

11 However, in contrast to the configuration shown in FIG. 1, the symmetric
12 configuration includes a probe client module 618 in the second network
13 device 604 that can perform similar probe sequence generation and analysis as
14 provided by the sustainable capacity probe module 610. The probe client
15 module 618 sends out probes from the second network device 604 to the first
16 network device 602, receives the corresponding responses and estimates the
17 sustainable capacity of the network from the second network device 604 to the
18 first network device 602, for circumstances when the sustainable capacities
19 available in the two directions differ.

20 FIG. 7 illustrates exemplary operations 700 for estimating sustained
21 capacity in a network. A transmission operation 702 transmits a first probe
22 sequence (e.g., containing ICMP ping packets) from a first network device to an IP
23 target. Each probe has a unique signature, such as a sequence number, that
24 distinguishes each probe packet from others in the sequence. The first probe
25 sequence provides a low load to the network (e.g., a negligible load to some load

1 less than the load provided by the second probe sequence). A reception
2 operation 704 receives the corresponding responses (e.g., ICMP ping ACK
3 packets) to the first probe sequence from the IP target. By collecting
4 corresponding probe-response pairs (e.g., based on the unique signatures), a
5 computation operation 706 determines the round-trip delay times for each pair in
6 the first probe sequence and then determines the average round-trip delay time \hat{D}_1
7 in the first probe sequence. The average round-trip delay time is one example of a
8 probe estimate. Other examples of a probe estimate may include without
9 limitation the median round-trip delay time, standard deviations in round-trip
10 delay times, round-trip delay time ranges, or other robust estimates.

11 A second transmission operation 708 transmits a second probe sequence
12 from the first network device to the IP target. The second probe sequence
13 provides a different load to the network (e.g., a load higher than that provided by
14 the first probe sequence). It should be understood, however, that the relative order
15 of the lower and higher loads may be changed. A reception operation 710 receives
16 the corresponding responses to the second probe sequence from the IP target. By
17 collecting corresponding probe-response pairs, a computation operation 712
18 determines the round-trip delay times for each pair in the second probe sequence
19 and then determines the average round-trip delay time \hat{D}_2 in the second probe
20 sequence.

21 An estimation operation 714 estimates the sustainable capacity available in
22 the network from the first network device to the IP target based on the average
23 round-trip delay times \hat{D}_1 and \hat{D}_2 , the difference in the loads δm , and a minimum
24 round-trip delay time D_{\min} associated with all of the probe-response pairs. In one
25 implementation, the estimate is computed as follows:

$$1 \quad \text{SusCap}_{est} = \delta m \frac{\hat{D}_2 - D_{\min}}{\hat{D}_2 - \hat{D}_1} \quad (1)$$

2 wherein SusCap_{est} represents the sustainable capacity that is available in the
 3 network channel from the first network device to the IP target. It should be
 4 understood that the same algorithm may also be employed in a symmetric
 5 configuration to determine the sustainable capacity that is available in the network
 6 from the IP target to the first network device. In another implementation, other
 7 statistics of the delay may be used instead of the average delays times. For
 8 example, \hat{D}_1 and \hat{D}_2 could include other probe estimates, such as the medians of
 9 the delays or trimmed means, among other metrics.

10 In yet another implementation, \hat{D}_1 and \hat{D}_2 may estimate the scale of the
 11 distributions. For example, \hat{D}_1 may be the standard deviation of the delays (or a
 12 robust alternative), with \hat{D}_2 being similarly defined, in which case, the estimate is
 13 computed as:

$$15 \quad \text{SusCap}_{est} = \delta m \frac{\hat{D}_2}{\hat{D}_2 - \hat{D}_1} \quad (2)$$

16 FIG. 8 illustrates alternative exemplary operations 800 for estimating
 17 sustained capacity in a network. Two heavily loaded probe sequences (i.e., at or
 18 near saturation of the network channel) are transmitted. However, either the probe
 19 packet sizes or the inter-probe gaps of the two sequences are different.

20 A delay trend of a given probe sequence may be estimated using the
 21 following algorithm:

$$23 \quad \text{DelayTrend}_i = \frac{1}{N-1} \sum_{n=2} (D_{i,n} - D_{i,n-1}) \approx \frac{1}{C} (p_i + B(\delta t_i)) - \delta t_i \quad (3)$$

1 wherein i represents the probe index (i.e., 1 or 2), N represents the number of
 2 probes packets in a sequence, $D_{i,n}$ represents the round-trip delay of a probe-
 3 response pair in the i -th sequence, $D_{i,n-1}$ represents the round-trip delay of the
 4 previous probe-response pair in the sequence, C represents the maximum network
 5 capacity, p_i represents the probe packet size of the i -th sequence, B represents the
 6 existing channel traffic (i.e., the background load), and δt_i represents the inter-
 7 probe gap of the i -th sequence. Therefore, the estimate of available sustainable
 8 capacity is given by $C-B$.

9 To determine the values of C and B , the delay trends of both sequences may
 10 be combined:

11 $DelayTrend_2 - DelayTrend_1$
 12 $\approx \left[\left[\frac{1}{C} (p_2 + B(\delta t_2)) - \delta t_2 \right] - \left[\frac{1}{C} (p_1 + B(\delta t_1)) - \delta t_1 \right] \right]$ (4)

14 If $\delta t_1 = \delta t_2$ and $p_1 \neq p_2$, and all are known values, then Equation (4) reduces to:

15 $C = \frac{p_2 - p_1}{DelayTrend_2 - DelayTrend_1}$ (5)

17 Thereafter, C can be inserted back into one of the $DelayTrend_i$ equations
 18 (e.g., $\frac{1}{C} (p_i + B(\delta t_i)) - \delta t_i$) to solve for B .

19 Alternatively, if $\delta t_1 \neq \delta t_2$ and $p_1 = p_2$, and all are known values, then the
 20 Equations (3) and (4) for $i=1$ and 2 can be solved to give C if some specified form
 21 of $B(\delta t)$ is assumed.

22 Accordingly, the implementation may be represented by the operations 800.
 23 A transmission operation 802 transmits a first probe sequence from a first network
 24 device to an IP target. The first probe sequence is characterized by a probe packet

1 size p_1 and an inter-probe gap δt_1 . Each probe has a unique signature, such as a
2 sequence number, that distinguishes each probe packet from others in the
3 sequence. A reception operation 804 receives the corresponding responses to the
4 first probe sequence from the IP target. By collecting corresponding probe-
5 response pairs (e.g., based on the unique signatures), a computation operation 806
6 determines the round-trip delay trend $DelayTrend_1$ for first probe sequence.

7 A second transmission operation 808 transmits a second probe sequence
8 from the first network device to the IP target. A reception operation 810 receives
9 the corresponding responses to the second probe sequence from the IP target. By
10 collecting corresponding probe-response pairs, a computation operation 812
11 determines the round-trip delay trend $DelayTrend_2$ for the second probe sequence.
12 A combination operation 814 combines the estimates of the delay trend to
13 determine the maximum network capacity C and the background load B . A
14 determining operation 816 subtracts the background load B from the maximum
15 network capacity C to determine the sustainable capacity in the network.

16 It should be understood that the same algorithm may also be employed in a
17 symmetric configuration to determine the sustainable capacity that is available in
18 the network from the IP target to the first network device.

19 Yet another implementation employs a windowing technique to determine a
20 saturation load on the network. A “window” represents a number of packets or a
21 number of bytes sent out from the server at any one time.

22 FIG. 9 illustrates an exemplary windowing technique for determining the
23 sustained capacity available in a network. A server 900 and a client 902 are
24 coupled by a network. Given a time period T , a slow start phase algorithm
25 incrementally increases the size of a window until packets are lost. For example,

1 at an early stage 904 of the slow start phase, a window 906 of seven packets (4
2 probe packets 908 and 3 response packets 910) is tested. At this stage, the
3 window 906 does not saturate the network. Therefore, the window 906 is
4 increased in size.

5 In a later stage 912 of the slow start phase, the window 906 is tested at 28
6 packets. At this window size, the network is full. After this window size, any
7 further increase in size would result in lost packets. Therefore, the slow start
8 phase stops at a maximum window size W_{max} equals 28.

9 In one implementation, the time period T is chosen to be the minimum
10 round trip delay time (RTT_{min}) detected by calibration (e.g., transmitting a
11 sequence of probes and taking the shortest round-trip delay time). A rough
12 estimate of the sustainable capacity may be computed using W_{max}/RTT_{min} , if the
13 window is in bytes

14
15 (or $\frac{W_{max}(\text{packet size})}{RTT_{min}}$, if the window size is in packets). (6)

16 In another implementation, a sustained transfer phase is used to flood the
17 network with the W_{max} number of packets for a predetermined time period T .
18 Then, the number of bytes per unit time (e.g., $\frac{W_{max}(\text{packet size})}{T}$) transmitted in the
19 window W_{max} during the predetermined time period T represents the sustainable
20 capacity of the network.

21 The exemplary hardware and operating environment of FIG. 10 for
22 implementing the invention includes a general purpose computing device in the
23 form of a computer 20, including a processing unit 21, a system memory 22, and a
24 system bus 23 that operatively couples various system components include the
25

1 system memory to the processing unit 21. There may be only one or there may be
2 more than one processing unit 21, such that the processor of computer 20
3 comprises a single central-processing unit (CPU), or a plurality of processing
4 units, commonly referred to as a parallel processing environment. The computer
5 20 may be a conventional computer, a distributed computer, or any other type of
6 computer; the invention is not so limited.

7 The system bus 23 may be any of several types of bus structures including a
8 memory bus or memory controller, a peripheral bus, a switched fabric, point-to-
9 point connections, and a local bus using any of a variety of bus architectures. The
10 system memory may also be referred to as simply the memory, and includes read
11 only memory (ROM) 24 and random access memory (RAM) 25. A basic
12 input/output system (BIOS) 26, containing the basic routines that help to transfer
13 information between elements within the computer 20, such as during start-up, is
14 stored in ROM 24. The computer 20 further includes a hard disk drive 27 for
15 reading from and writing to a hard disk, not shown, a magnetic disk drive 28 for
16 reading from or writing to a removable magnetic disk 29, and an optical disk drive
17 30 for reading from or writing to a removable optical disk 31 such as a CD ROM
18 or other optical media.

19 The hard disk drive 27, magnetic disk drive 28, and optical disk drive 30
20 are connected to the system bus 23 by a hard disk drive interface 32, a magnetic
21 disk drive interface 33, and an optical disk drive interface 34, respectively. The
22 drives and their associated computer-readable media provide nonvolatile storage
23 of computer-readable instructions, data structures, program modules and other
24 data for the computer 20. It should be appreciated by those skilled in the art that
25 any type of computer-readable media which can store data that is accessible by a

1 computer, such as magnetic cassettes, flash memory cards, digital video disks,
2 random access memories (RAMs), read only memories (ROMs), and the like, may
3 be used in the exemplary operating environment.

4 A number of program modules may be stored on the hard disk, magnetic
5 disk 29, optical disk 31, ROM 24, or RAM 25, including an operating system 35,
6 one or more application programs 36, other program modules 37, and program
7 data 38. A user may enter commands and information into the personal computer
8 20 through input devices such as a keyboard 40 and pointing device 42. Other
9 input devices (not shown) may include a microphone, joystick, game pad, satellite
10 dish, scanner, or the like. These and other input devices are often connected to the
11 processing unit 21 through a serial port interface 46 that is coupled to the system
12 bus, but may be connected by other interfaces, such as a parallel port, game port,
13 or a universal serial bus (USB). A monitor 47 or other type of display device is
14 also connected to the system bus 23 via an interface, such as a video adapter 48.
15 In addition to the monitor, computers typically include other peripheral output
16 devices (not shown), such as speakers and printers.

17 The computer 20 may operate in a networked environment using logical
18 connections to one or more remote computers, such as remote computer 49. These
19 logical connections are achieved by a communication device coupled to or a part
20 of the computer 20; the invention is not limited to a particular type of
21 communications device. The remote computer 49 may be another computer, a
22 server, a router, a network PC, a client, a peer device or other common network
23 node, and typically includes many or all of the elements described above relative
24 to the computer 20, although only a memory storage device 50 has been illustrated
25 in FIG. 10. The logical connections depicted in FIG. 10 include a local-area

1 network (LAN) 51 and a wide-area network (WAN) 52. Such networking
2 environments are commonplace in office networks, enterprise-wide computer
3 networks, intranets and the Internet, which are all types of networks.

4 When used in a LAN-networking environment, the computer 20 is
5 connected to the local network 51 through a network interface or adapter 53,
6 which is one type of communications device. When used in a WAN-networking
7 environment, the computer 20 typically includes a modem 54, a network adapter, a
8 type of communications device, or any other type of communications device for
9 establishing communications over the wide area network 52. The modem 54,
10 which may be internal or external, is connected to the system bus 23 via the serial
11 port interface 46. In a networked environment, program modules depicted relative
12 to the personal computer 20, or portions thereof, may be stored in the remote
13 memory storage device. It is appreciated that the network connections shown are
14 exemplary and other means of and communications devices for establishing a
15 communications link between the computers may be used.

16 In an exemplary implementation, probing modules, probe sequence
17 generators, probe performance analyzers, probe packet senders/receivers, and
18 other modules may be incorporated as part of the operating system 35, application
19 programs 36, or other program modules 37. Signatures, delay times, window
20 sizes, probe packets, response packets, and other data may be stored as program
21 data 38.

22 The embodiments of the invention described herein are implemented as
23 logical steps in one or more computer systems. The logical operations of the
24 present invention are implemented (1) as a sequence of processor-implemented
25 steps executing in one or more computer systems and (2) as interconnected

1 machine modules within one or more computer systems. The implementation is a
2 matter of choice, dependent on the performance requirements of the computer
3 system implementing the invention. Accordingly, the logical operations making
4 up the embodiments of the invention described herein are referred to variously as
5 operations, steps, objects, or modules.

6 The above specification, examples and data provide a complete description
7 of the structure and use of exemplary embodiments of the invention. Since many
8 embodiments of the invention can be made without departing from the spirit and
9 scope of the invention, the invention resides in the claims hereinafter appended.

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25